

1/16

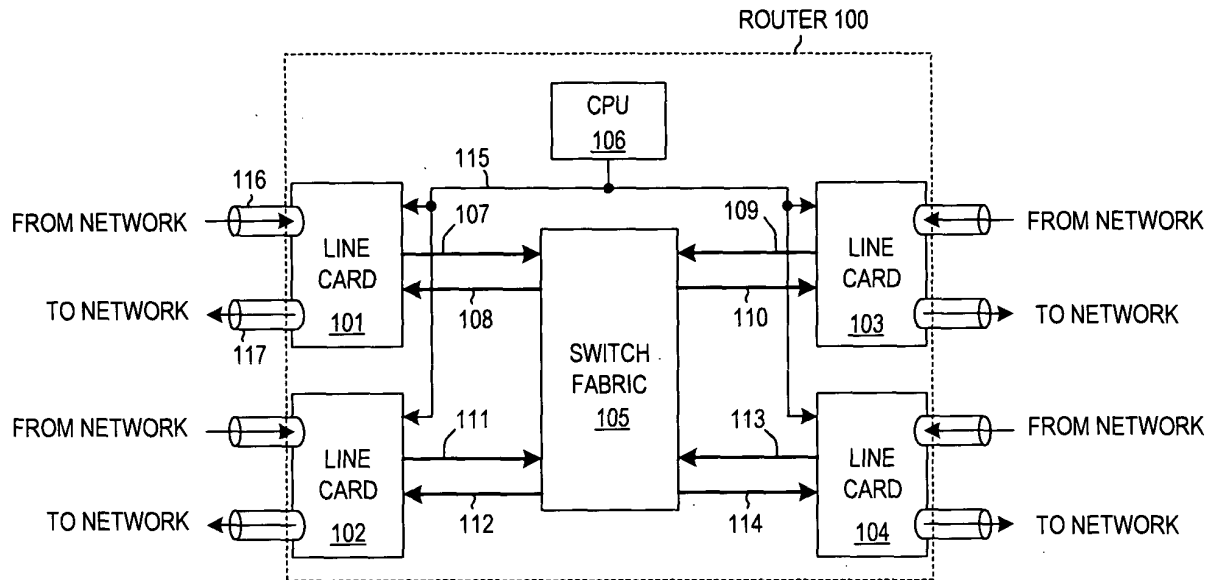


FIG. 1

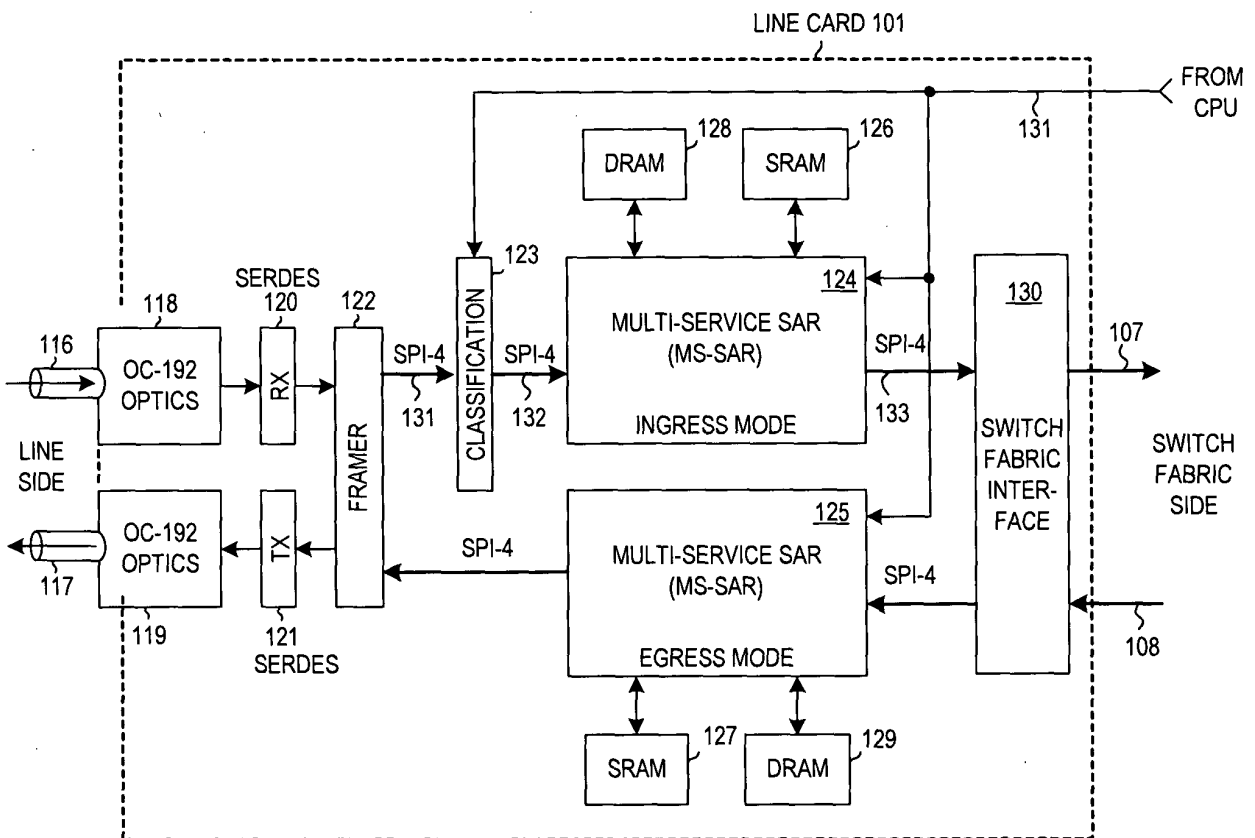


FIG. 2

2/16

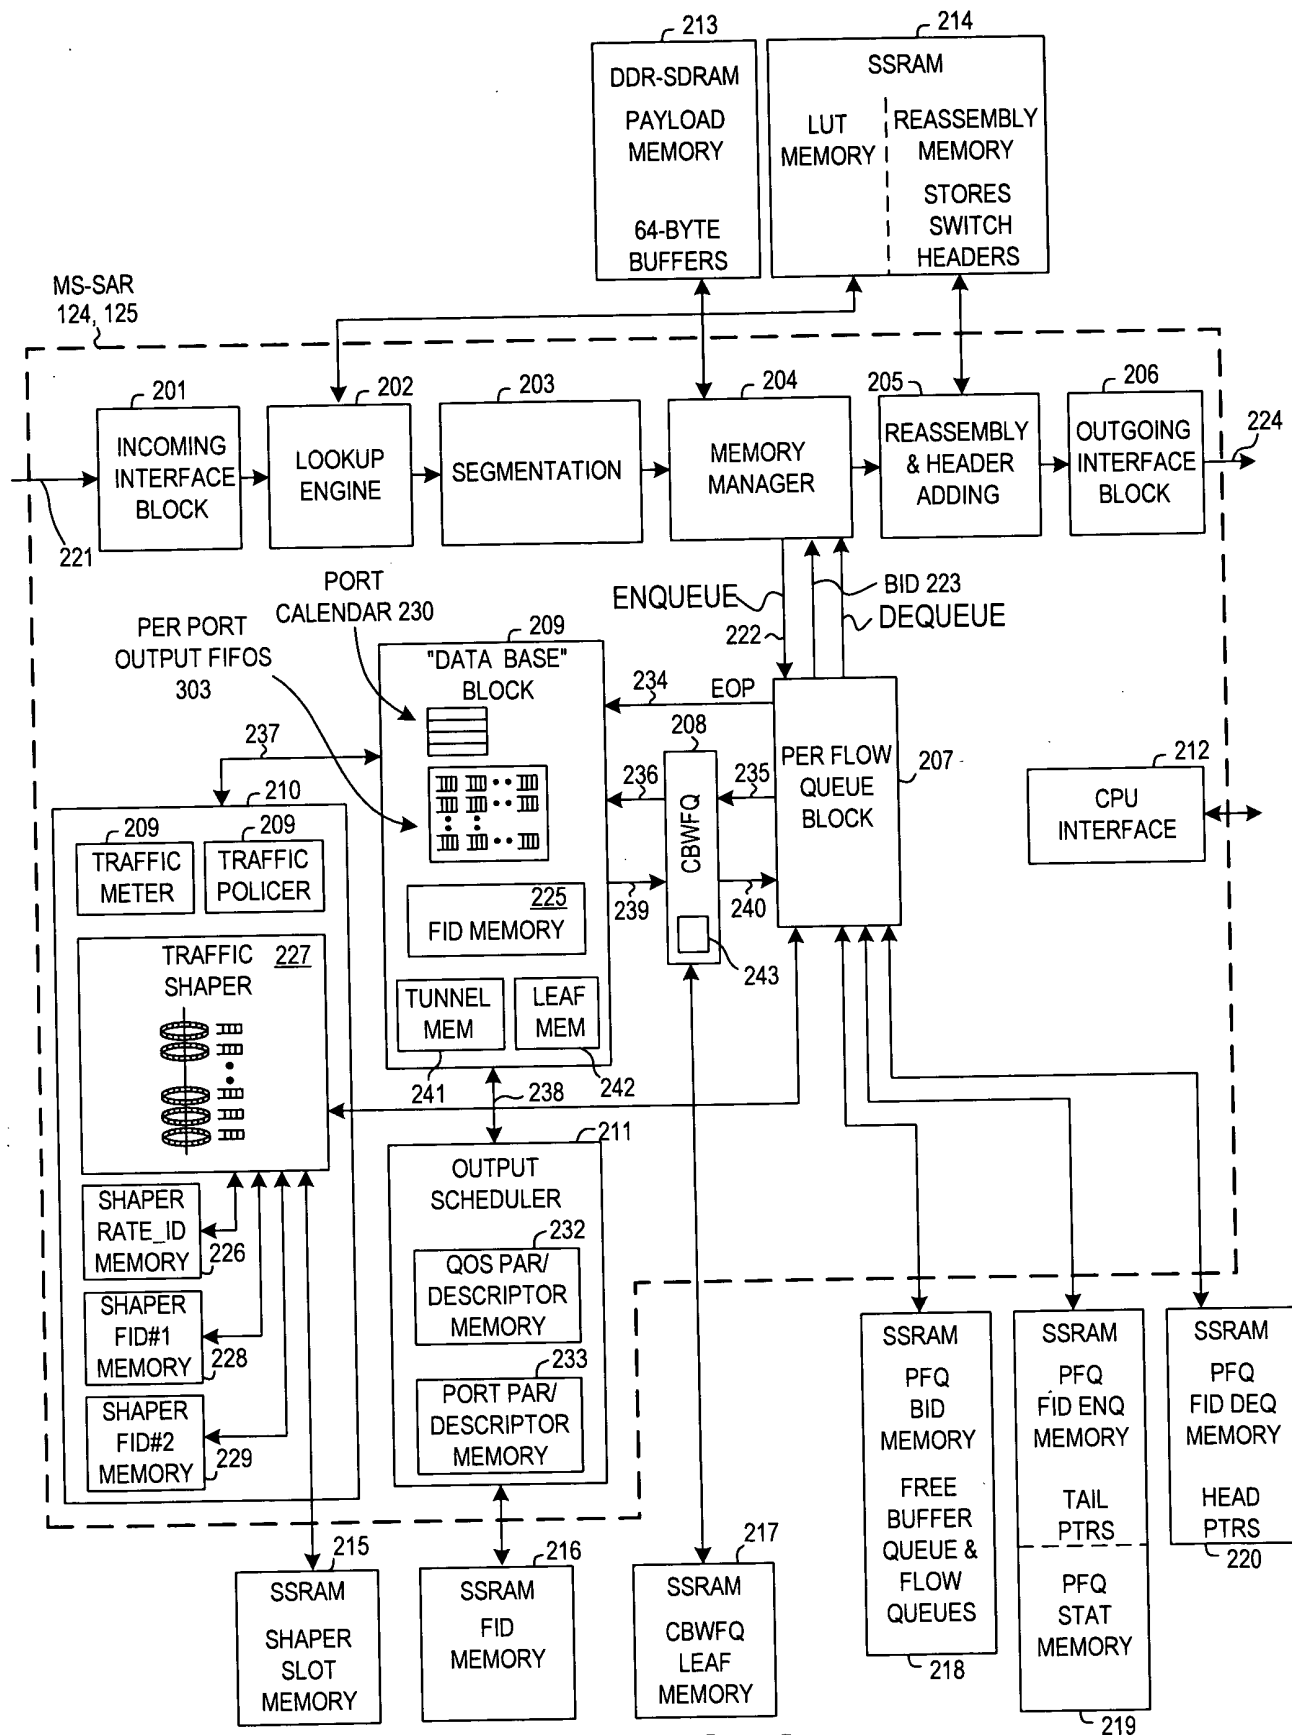


FIG. 3

3/16

FIELD NAME	NO BITS	RANGE	OWNED BY	DESCRIPTION
W_EMPTY	1	56	HARDWARE	IF SET, THE FLOW IS NOT IN THE SHAPER'S WHEEL.
W_F_EMPTY	1	55	HARDWARE	IF SET, THE FLOW IS TO BE REMOVED FROM THE WHEEL ON THE NEXT OUTPUT PHASE OF THE WHEEL WITHOUT LINKED TO THE OUTPUT PORT FIFOS.
P_EMPTY	1	54	HARDWARE	IF SET, THE FLOW IS NOT INSIDE THE SHAPER'S OUTPUT PORT FIFOS.
P_F_EMPTY	1	53	HARDWARE	IF SET, THE FLOW IS TO BE REMOVED FROM THE OUTPUT PORT FIFO ON THE NEXT OUTPUT PHASE OF THE DBS WITHOUT SENDING THE FLOW TO THE PFQ.
SCH_EMPTY	1	52	HARDWARE	IF SET, EN EMPTY INDICATION SHOULD BE SET TO THE SCHEDULER'S OUTPUT PHASE AND THE FLOWID SHOULD NOT BE TRANSFERRED TO THE PFQ.
CELL_CNT	12	51:40	HARDWARE	HOLDS THE AMOUNT OF SHAPED CELLS THAT ARE TO BE SENT TO THE PFQ. IF THIS VALUE IS NEGATIVE [11]=1, THEN DBS HAS TO ACCUMULATE CREDITS UNTIL IT BECOMES POSITIVE BEFORE SCHEDULING OUTPUT PHASE TO PFQ.
PARITY2	1	39	HARDWARE	PARITY FOR WORD 2.
PARITY1	1	38	HARDWARE	PARITY FOR WORD 1.
SHP_V	1	37	SOFTWARE	IF SET, THE FLOW BELONGS TO THE SHAPER.
RATE_ID	12	36:25	SOFTWARE	THE RATE_ID IS A VALUE THAT POINTS TO ONE OF 4K POSSIBLE SETS OF RATE VARIABLES (SHAPING PROFILES). THE RATE_ID POINTED TO WILL BE USED BY THE SHAPER IN SHAPING THE FID.
PORT	6	24:19	SOFTWARE	THE PORT WHICH THE FLOWID BELONGS TO.
SHP_CLASS	3	18:16	SOFTWARE	PRIORITY FOR THE SHAPED TRAFFIC. EACH SHAPED FID CAN EXIST IN ONLY ONE OF EIGHT STRICT PRIORITY CLASSES (QOS'S).
PARITY0	1	15	HARDWARE	PARITY FOR WORD 0.
FID_NEXT	15	14:0	HARDWARE	A POINTER TO NEXT FID IN THE OUTPUT PORT LINK.

DBS INTERNAL FID MEMORY

FIG. 4

4/16

FIELD NAME	NO BITS	OWNED BY	DESCRIPTION
EXPCDVT	5	SOFTWARE	EXPONENT OF CDVT.
MANCDVT	18	SOFTWARE	MANTISSA OF CDVT.
EXPKS	5	SOFTWARE	EXPONENT OF SUSTAIN RATE INCREMENT.
MANKS	18	SOFTWARE	MANTISSA OF SUSTAIN RATE INCREMENT.
EXPKP	5	SOFTWARE	EXPONENT OF PEAK RATE INCREMENT.
MANKP	18	SOFTWARE	MANTISSA OF PEAK RATE INCREMENT.
EXPINVKS	5	SOFTWARE	EXPONENT OF ONE OVER THE SUSTAIN RATE INCREMENT.
MANINVKS	18	SOFTWARE	MANTISSA OF ONE OVER THE SUSTAIN RATE INCREMENT.
THR	24	SOFTWARE	THRESHOLDS FOR SELECTING PEAK OR SUSTAIN RATE: [DIR_THR (1 BIT) ; EXP_THR (5 BITS) ;MAN_THR (18 BITS)]

INTERNAL SHAPER RATE_ID MEMORY

FIG. 5

FIELD NAME	NO BITS	OWNED BY	DESCRIPTION
RESERVED	3	HARDWARE	RESERVED.
FID_NEXT	15	HARDWARE	A POINTER TO THE NEXT FID FLOWID.

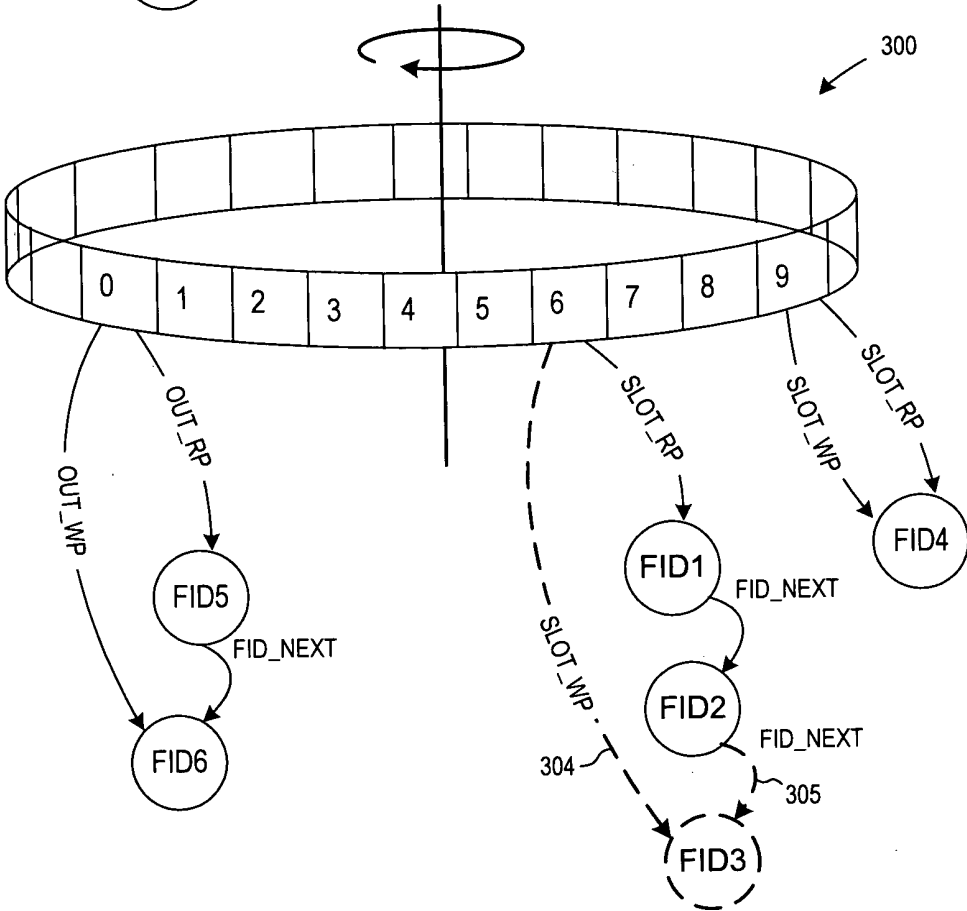
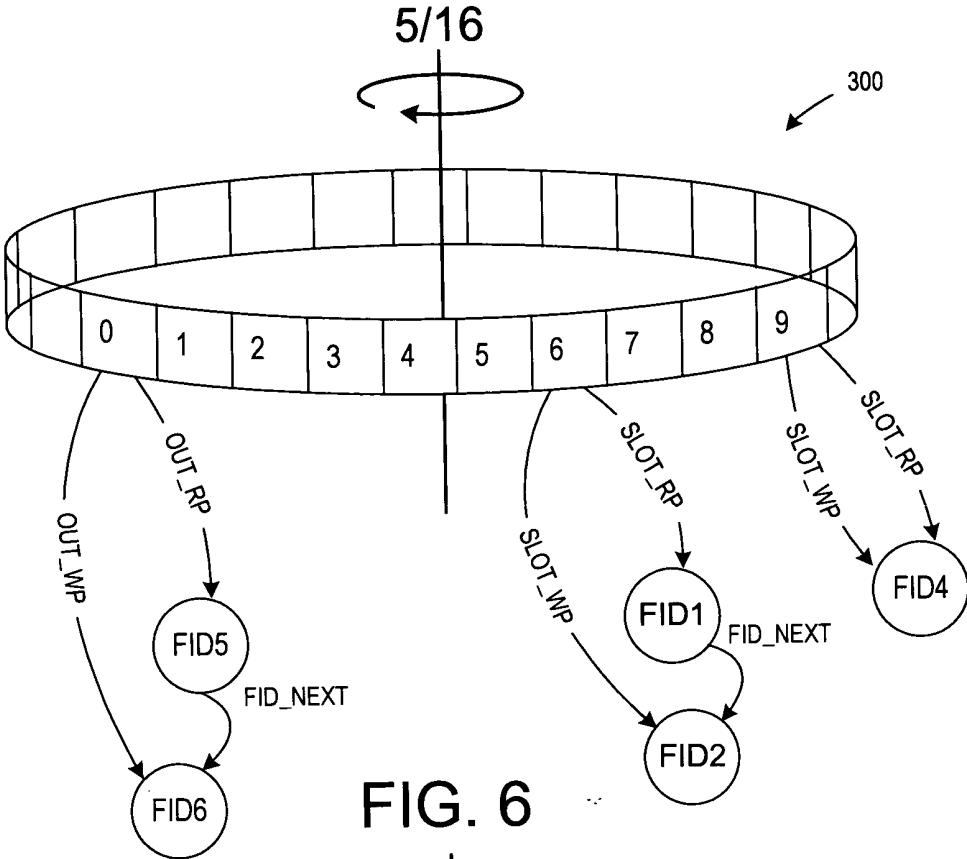
SHAPER INTERNAL FID MEM2

FIG. 8

FIELD NAME	NO BITS	OWNED BY	DESCRIPTION
SLOT_RP	16	HARDWARE	READ POINTER OF THE FID LINKED LIST.
SLOT_WP	16	HARDWARE	WRITE POINTER OF THE FID LINKED LIST.

SHAPER EXTERNAL SLOT MEMORY

FIG. 10



6/16

FIELD NAME	NO BITS	OWNED BY	DESCRIPTION
PLC_V_1	1	SOFTWARE	ACTIVE LOW. IF 0, THIS FID IS ASSIGNED FOR POLICING. NO SHAPING IS ALLOWED. IF 1, THIS FID IS NOT FOR POLICING. THE FID CAN BE USED FOR ANY OTHER FUNCTIONS INCLUDING SHAPING.
RESIDUE	16	HARDWARE	RESIDUE FROM INPUT/OUTPUT PHASE CALCULATION.
BG_STT	2	HARDWARE	FOR DUAL LEAKY BUCKET. SET TO 11 DURING SETUP CONNECTION COMMAND USED TO MARK THE FLOWID FOR TIME COUNT WRAPPING AROUND. USED FOR P_S/CREDIT CALCULATION IN DUAL LEAKY BUCKET ONLY.
TIME_MSB	1	HARDWARE	BIT[18] (THE 19TH BIT) OF FREE RUNNING TIME COUNTER. USED FOR P_S/CREDIT CALCULATION IN DUAL LEAKY BUCKET ONLY.
ARRIVAL_TIME	18	HARDWARE	TIMESTAMP OF WHEN THE FID ARRIVED INTO THE SHAPER WHEEL OUTPUT FIFO.
ADDITION FOR DUAL LEAKY BUCKET:			
LAST_TIME	20	HARDWARE	TIMESTAMP OF THE LAST TIME THIS FID WAS SENT FROM THE SHAPER TO THE DATA BASE BLOCK. THIS WAS THE TIME P_S/ CREDIT WAS UPDATED LAST. USED FOR P_S/CREDIT CALCULATION IN DUAL LEAKY BUCKET ONLY.
SHP_CREDIT	24	HARDWARE	AMOUNT OF CREDIT ACCUMULATED FOR THE FLOWID. USED FOR P_S/CREDIT CALCULATION IN DUAL LEAKY BUCKET.
PEAK_SUSTAIN	1	HARDWARE	IF SET, USE PEAK RATE TO CALCULATE FUTURE TIME SLOT DURING OUTPUT PHASE. IF RESET, USE SUSTAINED RATE TO CALCULATE FUTURE SLOT DURING OUTPUT PHASE.

WRITE TO FIELD	MEMORY BITS
PLC_V_1	[0]
RESIDUE [15:2]	[14:1]
BG_STT	[16:15]
TIME_MSB	[17]
ARRIVAL_TIME	[35:18]
LAST_TIME	[55:36]
CREDIT_MAN [17:8]	[65:56] [69:66]
CREDIT_EXP [3:0]	[70]
CREDIT_DIR	[71]
P_S	
RESIDUE [1:0]	[73:72]
CREDIT_MAN [7:0]	[81:74]
CREDIT_EXP [4]	[82]

SHAPER INTERNAL FID MEM1

FIG. 7

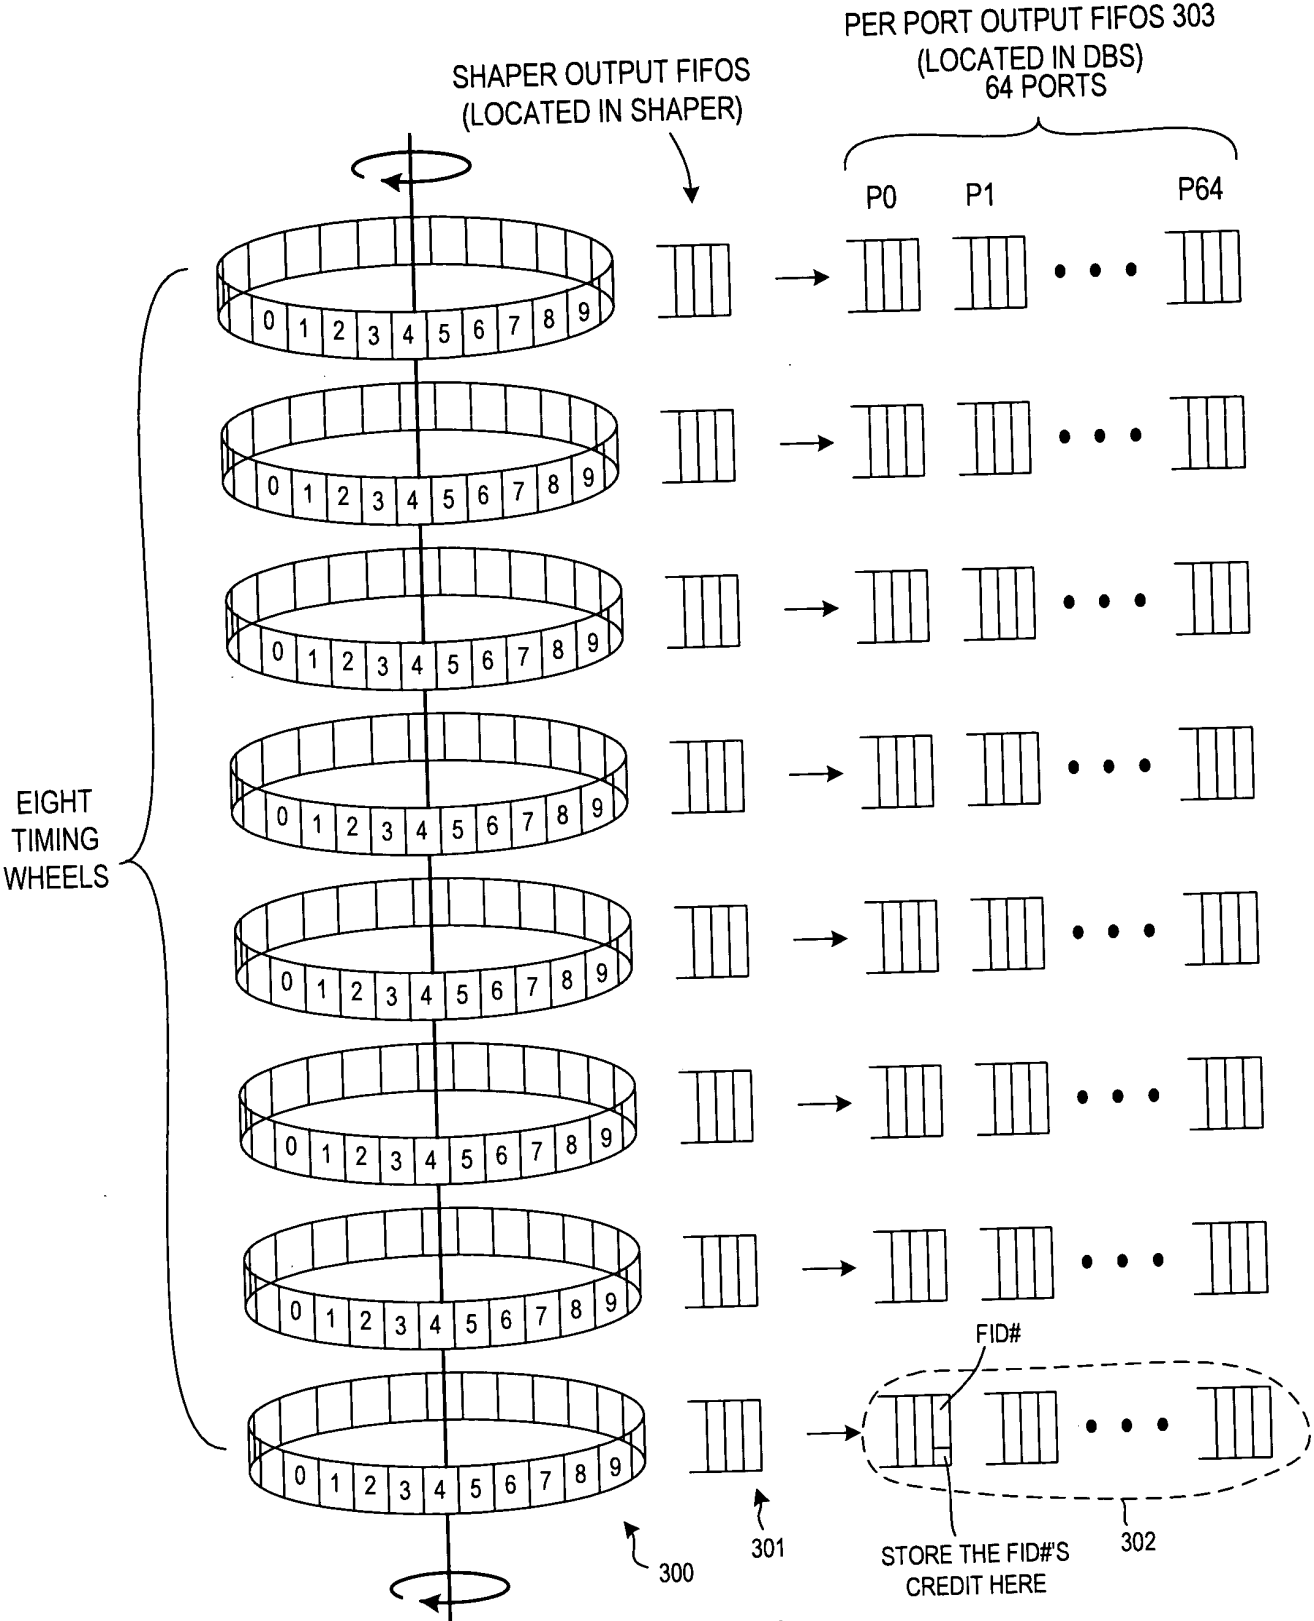


FIG. 11

8/16

FIELD NAME	LOCATION	NO BITS	OWNED BY	DESCRIPTION
F_NEXT	35:16	20	SCHEDULER	THE NEXT FID, WHICH IS LINKED TO THE PREVIOUS FID. THE ADDRESS TO THE MEMORY IS THE CURRENT FID.
EMPTY	15	1	SCHEDULER	IF EMPTY = 1, THEN THE FID IS NEW AND HAS TO BE LINKED BY THE INPUT PHASE. THE OUTPUT PHASE CAN ALSO MAKE EMPTY=1 IF IT RECEIVES THE ASSERTED EMPTY BIT FROM DBS.
SCH_V	14	1	SOFTWARE	IF SET, THE FID BELONGS TO THE SCHEDULER.
PARITY	13	1	HARDWARE	ODD PARITY BIT FOR DATA.
RESERVED	12	1		
QOS_ADDRESS	11:0	12	SOFTWARE	QOS ADDRESS TO THE QOS MEMORIES [PORT, QOS]. BITS [11:3] ARE THE PORT NUMBER. BITS [2:0] ARE THE QOS. THERE ARE ELEVEN BITS FOR PORTS AND THUS SUPPORT FOR 2K PORTS. ACTUAL NUMBER OF PORTS SUPPORTED IS 1K (VIRTUAL) OR 64 LOGICAL OR COMBINATION OF BOTH. THESE PORT NUMBERS HAVE DIFFERENT MEANINGS DEPENDING ON MODE. MODE 0: ONLY 64 PORTS ARE USED OUT OF POSSIBLE 2K PORTS. ONLY VALUES [63:0] ARE VALID. MODE 1: PORTS [0X3FF:0X404] ARE USED TO MAP LOGICAL PORTS [63:4]. MODE 2: PORTS [0X3FF:0] ARE USED FOR VIRTUAL PORT MAPPING TO ONE LOGICAL PORT[0] AND PORTS [0X43F:0X401] ARE USED TO MAP LOGICAL PORTS [63:1]. IF QOS ADDRESS IS 0X3FFF, THEN IT BELONGS TO THE CPU PORT.

SCHEDULER EXTERNAL FID MEMORY

FIG. 12

9/16

FIELD NAME	LOCATION	NO BITS	OWNED BY	DESCRIPTION
PARITY	66	1		EVEN PARITY FOR DATA [65:0].
Q_WEIGHT_MF	65:41	25	SCHEDULER	WEIGHT COUNT DOWN VALUE THAT STORES THE AMOUNT OF WEIGHT REMAINING FOR THE QOS. SINCE THIS FIELD CAN BECOME NEGATIVE, THERE IS AN EXTRA BIT TO SAVE THE SIGN.
F_RP	40:21	20	SCHEDULER	THE HEAD OF THE FID LINKED LIST FOR THE QOS.
F_WP	20:1	20	SCHEDULER	THE TAIL OF THE FID LINKED LIST FOR THE QOS.
F_EMPTY	0	1	SCHEDULER	AN INDICATION OF AN EMPTY LINKED LIST.

SCHEDULER INTERNAL QOS DESCRIPTOR MEMORY

FIG. 13

FIELD NAME	LOCATION	NO BITS	OWNED BY	DESCRIPTION
PARITY	8	1		PARITY FOR BITS [7:0].
Q_WEIGHT	7:0	8	SOFTWARE	WEIGHT OF THE QOS FOR THE WEIGHTED ROUND ROBIN ALGORITHM.

SCHEDULER INTERNAL QOS PARAMETER MEMORY

FIG. 14

FIELD NAME	LOCATION	NO BITS	OWNED BY	DESCRIPTION
PARITY	11	1		EVEN PARITY FOR BITS [10:0].
PRIORITY	10:3	8	SOFTWARE	EACH BIT IN THIS REGISTER INDICATES WHETHER THE ASSOCIATED QOS FOR THE OUTPUT PORT IS STRICT PRIORITY OR IS WRR. A BIT SET TO "1" INDICATES STRICT PRIORITY. A BIT SET TO "0" INDICATES WRR. THE RIGHT-MOST BIT IS QOS 0. STRICT AND WRR BITS CANNOT BE INTERLEAVED. QOS 7 WILL ALWAYS BE BEST EFFORTS IRRESPECTIVE OF WHAT IS SET FOR QOS 7.
FACTOR	2:0	3	SOFTWARE	IN ORDER TO MAKE THE TRAFFIC LESS BURSTY, EACH QOS WEIGHT WILL BE DIVIDED BY 2^{FACTOR} . THAT MEANS THAT THERE WILL BE SOME ROTATION OF QOS'S IN THE ACTIVE LINKED LIST BEFORE MOVING THEM TO THE WAITING LIST.

SCHEDULER INTERNAL PORT PARAMETER MEMORY

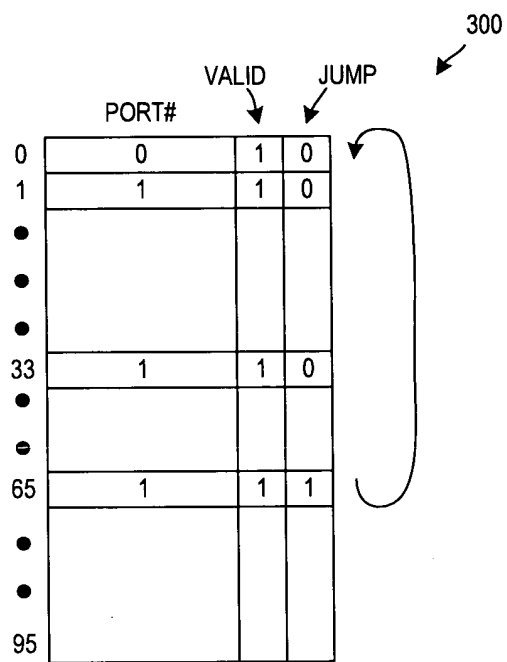
FIG. 15

10/16

FIELD NAME	LOCATION	NO BITS	OWNED BY	DESCRIPTION
PARITY1	48	1	SCHEDULER	PARITY FOR BITS [47:9].
PREV_QOS	47:45	3	SCHEDULER	THE QOS THAT WAS SERVICED LAST TO BE SERVICED FOR THIS PORT.
PREV_QOS_V	44	1	SCHEDULER	IF SET, PREV_QOS HAS TO BE SCHEDULED WITH NO CONNECTION TO PRIORITY SINCE IN THE MIDDLE OF A PACKET. IF RESET, CHOOSE A QOS ACCORDING TO PRIORITY AND WRR.
ACTIVE_PTR	43:41	3	SCHEDULER	THIS POINTER POINTS TO CURRENT QOS TO BE SERVICED IN ROUND ROBIN FASHION. WHEN Q_WEIGHT_M IS EITHER ZERO OR NEGATIVE AND EOP IS RECEIVED FOR THE CURRENT FID, THEN THIS POINTER MOVES TO THE NEXT QOS IN ROUND ROBIN. THIS POINTER ALSO MOVES WHEN EMPTY (NO MORE PACKET PENDING FOR THIS FID) IS RECIEVED FROM DBS AND F_RP AND F_WP ARE THE SAME, IRRESPECTIVE OF Q_WEIGHT_M AS THE LINKED LIST BECOMES EMPTY FOR THIS QOS. NOTE THAT THIS POINTER MOVES AROUND ONLY ON ROUND ROBIN QOS'S. WHEN A STRICT PRIORITY QOS IS SERVICED, THIS POINTER IS NOT ALTERED.
Q_WEIGHT_MF	40:17	24	SCHEDULER	A COUNT DOWN COUNTER OF WEIGHT/FACTOR FOR THE CURRENT QOS. WHEN THIS COUNT BECOMES ZERO AND EOP ARRIVES, NEXT QOS AVAILABLE WILL BE SERVICED AND A NEW QOS WEIGHT/PORT FACTOR IS LOADED INTO THIS COUNTER.
QA_EMPTY	16:9	8	SCHEDULER	ONE BIT INDICATION (PER QOS) IF THE QOS IS EMPTY (OR IF SOME FIDS ARE LINKED TO IT). THIS ACTIVE EMPTY LIST IS USED FOR THE PRIORITY SCHEME. WHEN ANY OF THE QOS WEIGHT BECOMES ZERO OR NEGATIVE, THEN IT IS MOVED TO THE QW_EMPTY LIST. QW_EMPTY IS LOADED TO THIS LIST WHEN ALL ROUND ROBIN QOS'S WEIGHT BECOME ZERO OR NEGATIVE. IF ANY STRICT PRIORITY QOS IS PENDING IN QW_EMPTY LIST, THEN IT IS SERVICED FIRST BEFORE RESUMING ROUND ROBIN OPERATION.
PARITY0	8	1	SCHEDULER	PARITY FOR FIRST EIGHT BITS.
QW_EMPTY	7:0	8	SCHEDULER	ONE BIT INDICATION (PER QOS) IF THE QOS IS WAITING EMPTY. THERE IS ONE BIT PER EACH OF THE EIGHT QOS'S. SCHEDULER INPUT PHASE UPDATES THIS LIST TO INDICATE THAT THERE IS TRAFFIC FOR A QOS. ALSO NON_EMPTY QOS'S ARE MOVED FROM ACTIVE LIST TO THIS LIST WHEN THERE IS EOP AND WEIGHT IS DECREMENTED TO ZERO OF NEGATIVE.

SCHEDULER INTERNAL PORT DESCRIPTOR MEMORY

FIG. 16



PORT CALENDAR IN DATA BASE BLOCK

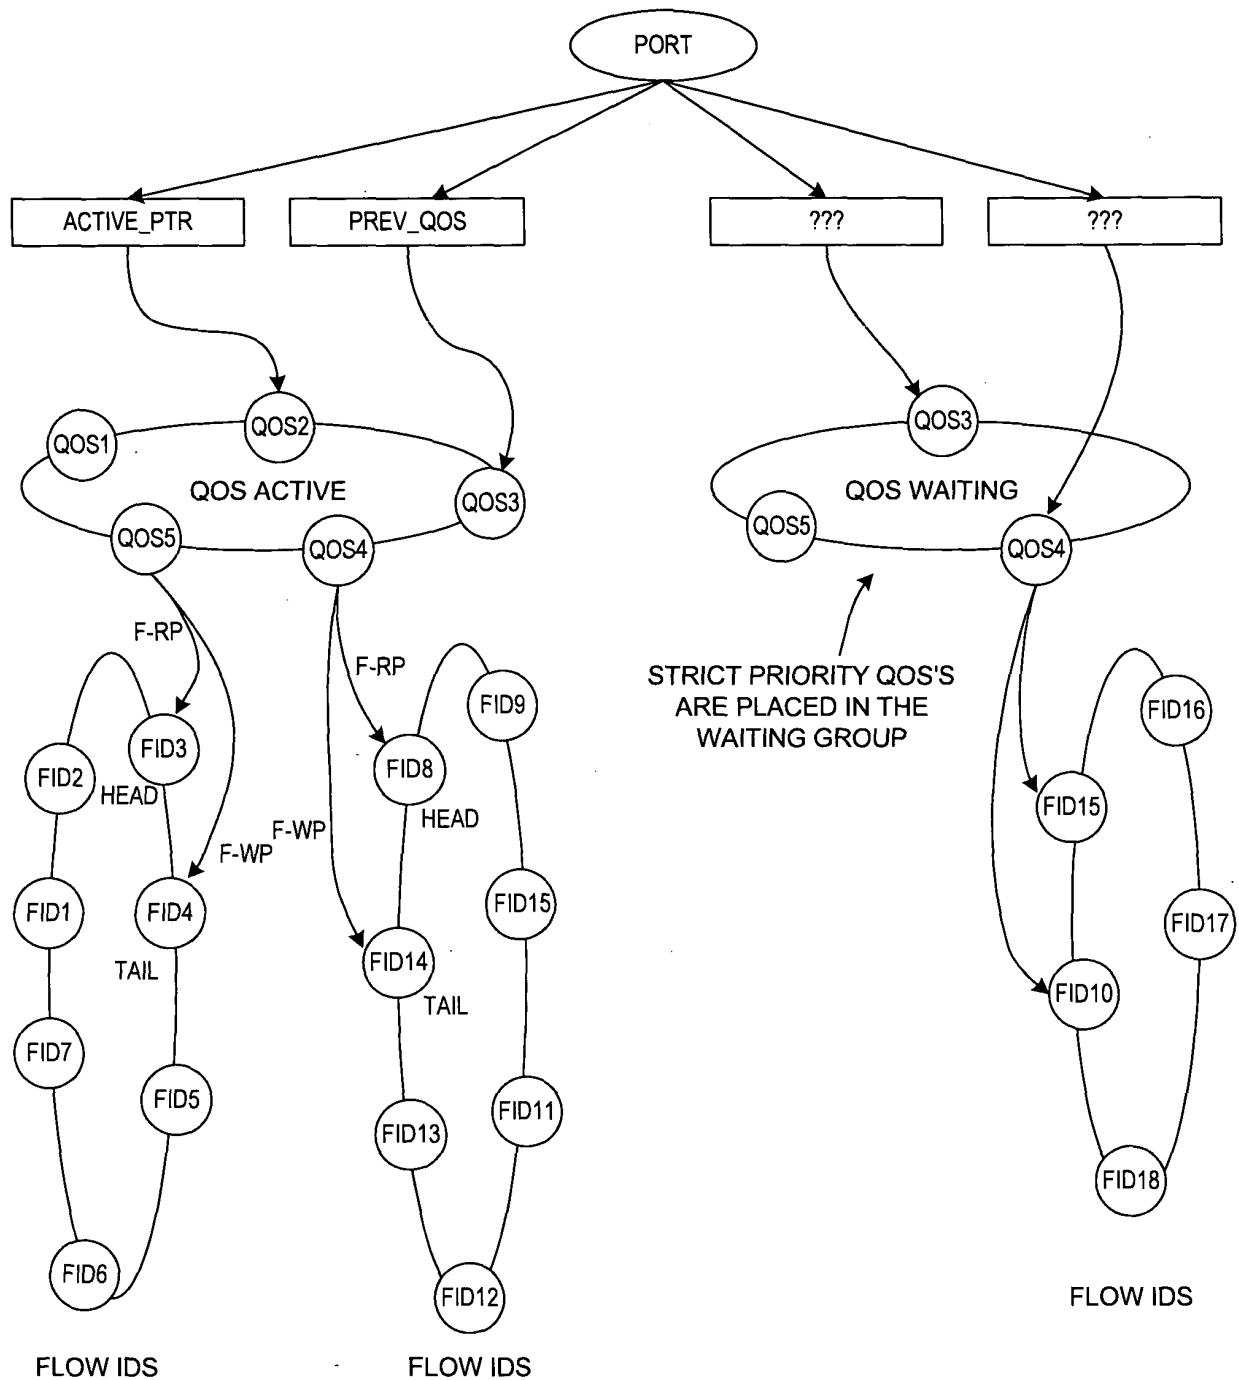
FIG. 17

FIELD NAME	NO BITS	RANGE	OWNED BY	DESCRIPTION
PORTID1	6	[5:0]	SOFTWARE	THE VALUE OF THE OUTPUT PORT.
VALID1	1	[6]	SOFTWARE	IF SET, THE PORTID1 IS VALID. IF RESET, THE PORTID1 IS NOT VALID.
JUMP1	1	[7]	SOFTWARE	IF SET, THE ADDRESS TO THIS MEMORY SHOULD BE RESET IN THE NEXT POSITIVE EDGE OF THE CLOCK. IF RESET, THE ADDRESS TO THIS MEMORY SHOULD BE INCREMENTED BY 1 IN THE NEXT POSITIVE EDGE OF THE CLOCK.
PARITY1	1	[8]	HARDWARE	PARITY BIT 1.
PORTID2	6	[14:9]	SOFTWARE	THE VALUE OF THE OUTPUT PORT.
VALID2	1	[15]	SOFTWARE	IF SET, THE PORTID2 IS VALID. IF RESET, THE PORTID2 IS NOT VALID.
JUMP2	1	[16]	SOFTWARE	IF SET, THE ADDRESS TO THIS MEMORY SHOULD BE RESET IN THE NEXT POSITIVE EDGE OF THE CLOCK. IF RESET, THE ADDRESS TO THIS MEMORY SHOULD BE INCREMENTED BY 1 IN THE NEXT POSITIVE EDGE OF THE CLOCK.
PARITY2	1	[17]	HARDWARE	PARITY BIT 2.

PORT CALENDAR MEMORY

FIG. 18

12/16



WEIGHTED ROUND ROBIN SCHEDULING

FIG. 19

FIELD NAME	NO BITS	RANGE	OWNED BY	DESCRIPTION
PARITY0	1	[30]	HARDWARE	PARITY BIT.
TUNNEL_VALID	1	[29]	SOFTWARE	IF SET, THIS FID IS USED AS A TUNNEL. IF RESET, THIS FID IS USED AS A REGULAR FID. USED ON THE OUTPUT PHASE TO DECIDE IF THE TUNNEL IS SENT TO THE PFQ OF THE LEAF.
LEAF_RP	14	[28:15]	HARDWARE	READ POINTER TO THE LIST OF NON-EMPTY LEAFS.
LEAF_WP	14	[14:1]	HARDWARE	WRITE POINTER TO THE LIST OF NON-EMPTY LEAFS.
TUNNEL_EMPTY	1	[0]	HARDWARE	IF SET, THE NON-EMPTY LEAFS LIST IS EMPTY.

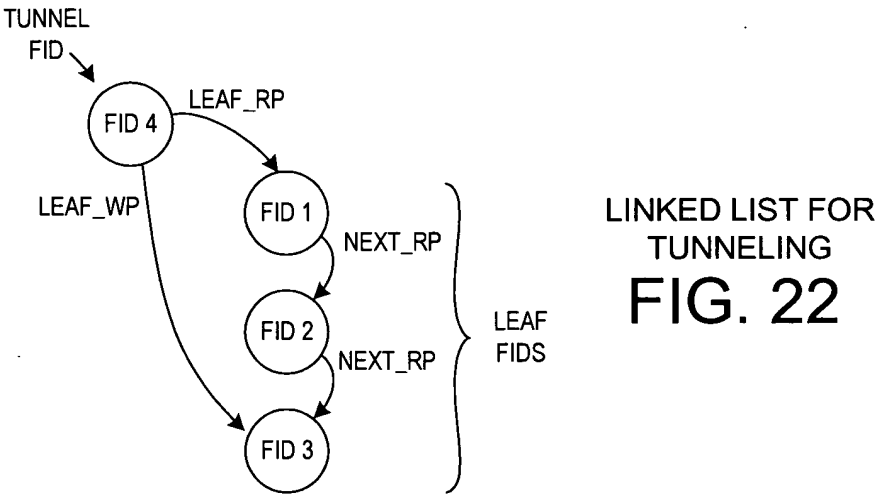
TUNNEL MEMORY INSIDE DATA BASE BLOCK

FIG. 20

FIELD NAME	NO BITS	RANGE	OWNED BY	DESCRIPTION
LEAF_VALID	1	[27]	SOFTWARE	IF SET, THIS FID IS USED AS A LEAF INSIDE A TUNNEL. IF RESET, THIS FID IS USED AS A REGULAR FID. USED ON THE INPUT PHASE TO DECIDE IF THE LEAF OR THE TUNNEL IS SENT TO THE SHAPER.
TUNNEL_PTR	10	[26:17]	SOFTWARE	A POINTER TO THE TUNNEL THIS LEAF BELONGS TO.
LEAF_EMPTY	1	[16]	HARDWARE	IF SET, THE LEAD FID IS EMPTY.
PARITY1	1	15	HARDWARE	PARITY FOR WORD 1.
NEXT_LEAF	14	[14:1]	HARDWARE	A POINTER TO THE NEXT NON-EMPTY LEAF (LINK LIST)
PARITY0	1	0	HARDWARE	PARITY FOR WORD 1.

LEAF MEMORY INSIDE DATA BASE BLOCK

FIG. 21



14/16

FIELD NAME	LOCATION	NO BITS	OWNED BY	DESCRIPTION
LOCATION 0				
WEIGHT	[7:0]	8	SOFTWARE	WEIGHT OF THE IFID FOR THE WFQ ALGORITHM.
ROOT_ADDRESS	[21:8]	14	SOFTWARE	A POINTER TO THE ROOT FID.
RESERVED	[25:22]	4		
LEAF_VALID	LOCATION	1	SOFTWARE	IF SET, THIS FID IS A LEAF (SHOULD BE MERGED). IF RESET, THIS FID IS A REGULAR FLOW.
HIGH_PRIORITY	[27]	1	SOFTWARE	A HIGH PRIORITY INDICATION FOR THE MDRR AND PRIORITY AND DRR ALGORITHMS.
PARITY0	[28]	1	HARDWARE	ODD PARITY BIT IS WRITTEN AND CHECKED FOR ALL MEMORY ACCESSSES.
RESERVED	[35:29]	7		
LOCATION 1				
WEIGHT_CNT	[15:0]	16	HARDWARE	CURRENT COUNT OF CELLS THAT WERE SENT ALREADY.
LEAF_EMPTY	[16]	1	HARDWARE	AN INDICATION IF THE LEAF IS EMPTY.
PARITY1	[17]	1	HARDWARE	PARITY BIT FOR BITS [16:0].
LEAF_NEXT	[34:18]	17	HARDWARE	POINTER TO THE NEXT IFID IN THE GROUP.
PARITY2	[35]	1	HARDWARE	PARITY BIT FOR LEAF NEXT.

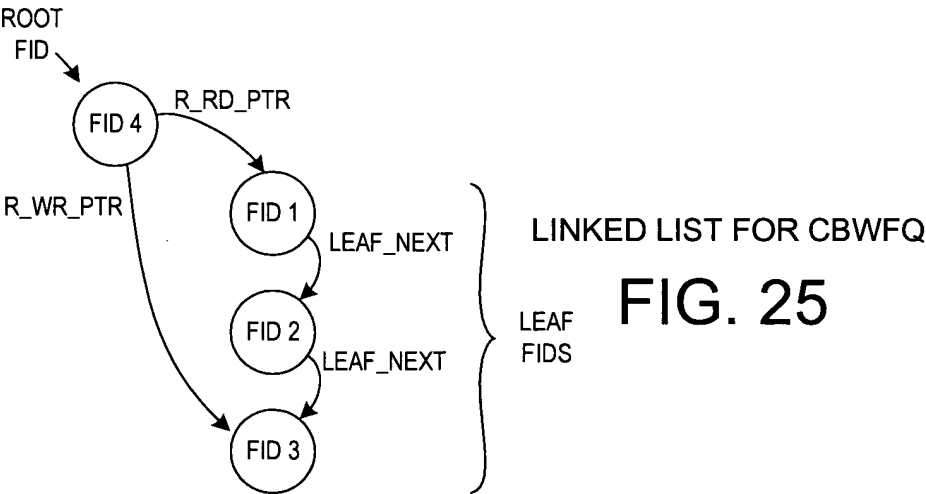
EXTERNAL CBWFQ LEAF DESCRIPTOR MEMORY

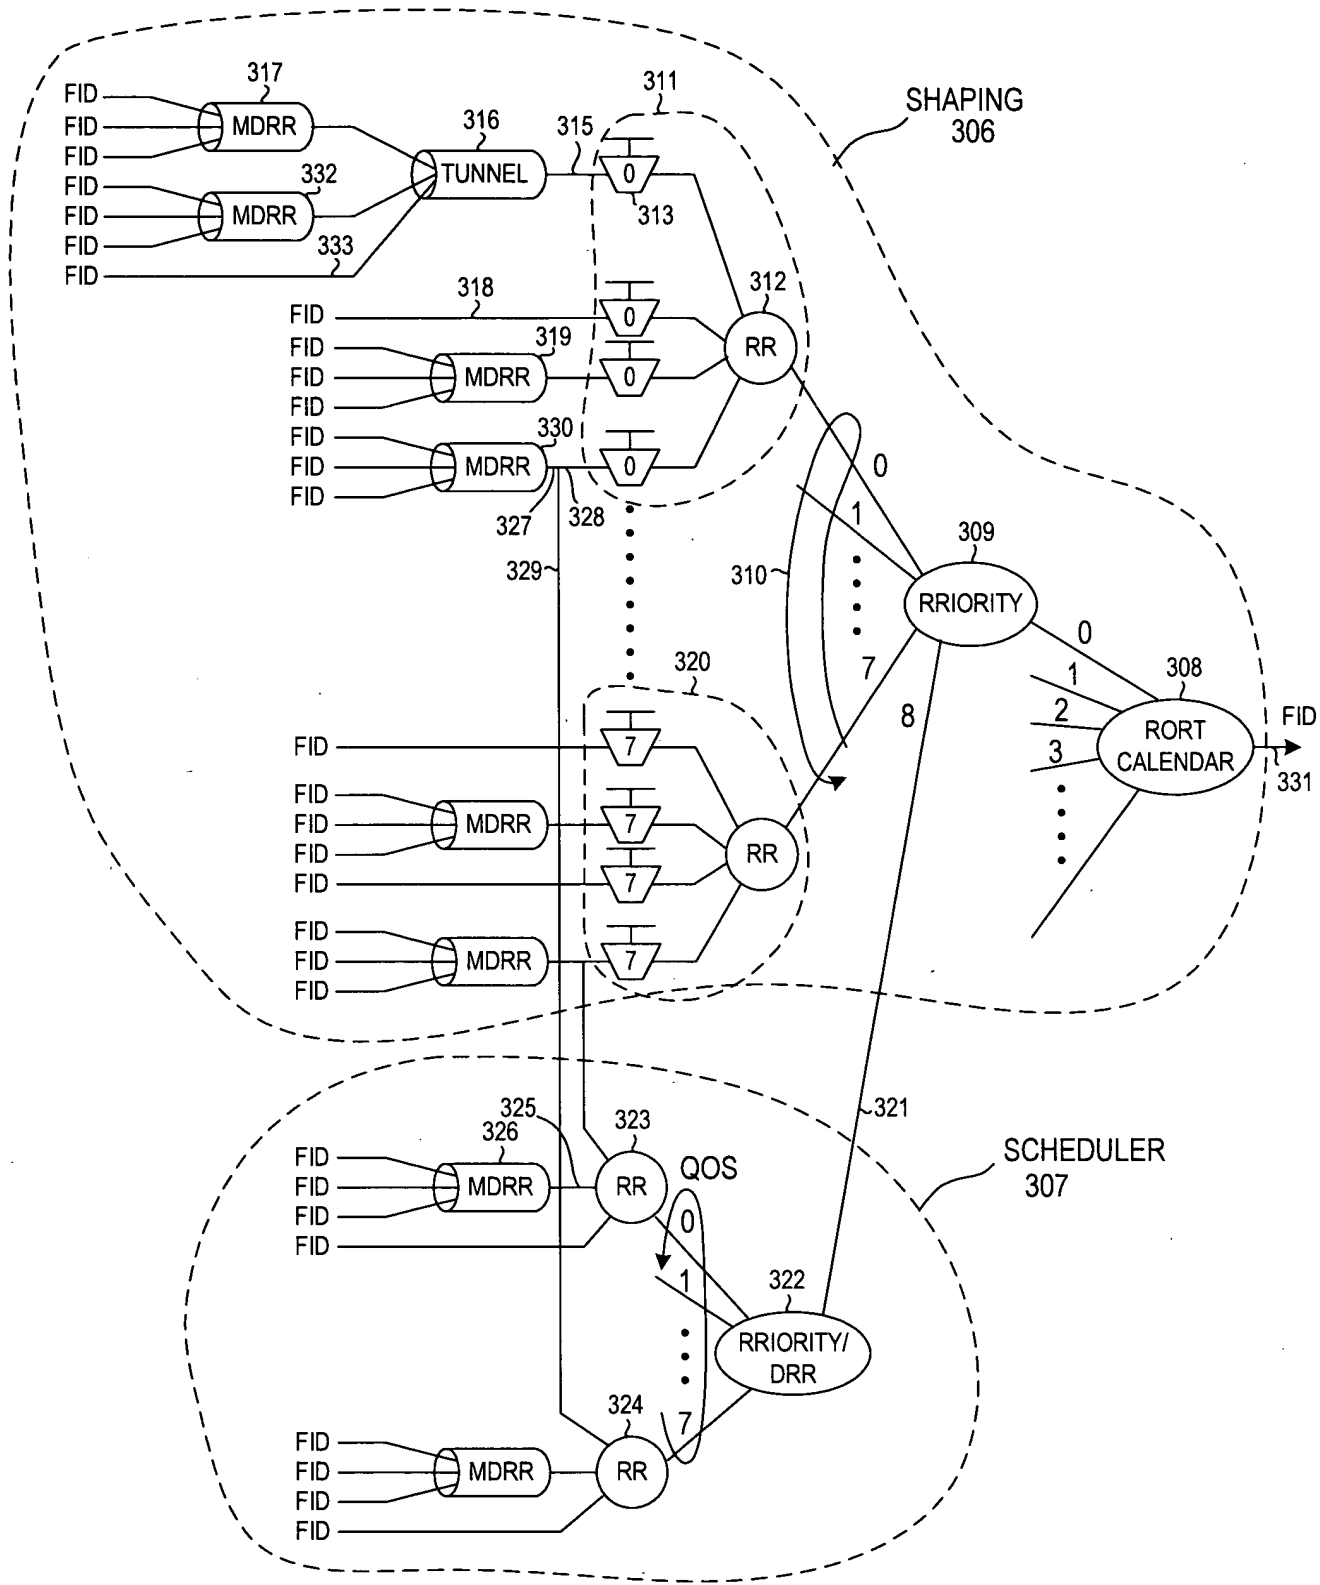
FIG. 23

FIELD NAME	LOCATION	NO BITS	OWNED BY	DESCRIPTION
ROOT_VALID	[0]	1	SOFTWARE	IF SET, THIS FID IS A ROOT. IF RESET, THIS FID IS A REGULAR FLOW.
R_MODE	[3:1]	3	SOFTWARE	MODE OF OPERATION PER VC: 000 - RR 001 - DRR 010 - MDRR 011 - PRIORITY AND DRR 100 - PRIORITY AND DRR AND DISCARD 101 - 111 - RESERVED.
HIGH_FID	[20:4]	17	SOFTWARE	A POINTER TO THE HIGH PRIORITY LEAF FID.
R_RD_PTR	[37:21]	17	HARDWARE	A POINTER TO THE HEAD OF THE ACTIVE LEAF'S LIST.
R_WR_PTR	[54:38]	17	HARDWARE	A POINTER TO THE TAIL OF THE ACTIVE LEAF'S LIST.
R_EMPTY	[55]	1	HARDWARE	AN EMPTY INDICATION FOR THE ROOT.
HIGH_EMPTY	[56]	1	HARDWARE	AN EMPTY INDICATION TO THE HIGH PRIORITY FID.
HIGH_RR	[57]	1	HARDWARE	WHEN THIS BIT IS SET, HIGH FID IS BEING SERVICED. WHEN THIS BIT IS LOW, PRIORITY FID IS BEING SERVICED.
WAIT_EOP	[58]	1	HARDWARE	WHEN THIS BIT IS SET, FID SERVICED LAST DID NOT RECEIVE EOP, HENCE SERVICE SAME FID WITHOUT CONSIDERING ANY PRIORITY.
NEG_WEIGHT	[59]	1	HARDWARE	THIS BIT IS SET BY HARDWARE WHEN THE NEXT LINK TO BE SERVICED HAS STILL NEGATIVE WEIGHT AFTER ADDING NEW QUOTA*WEIGHT. THIS BIT HAS TO BE SET IN ORDER FOR THE OUTPUT PHASE TO NOT TO SERVE THE CURRENT FID IF THIS BIT IS SET.
PARITY_1	[60]	1	HARDWARE	PARITY FOR BITS [37:0].
PARITY_H	[61]	1	HARDWARE	PARITY FOR BITS [58:38].

INTERNAL ROOT (VC) DESCRIPTOR MEMORY

FIG. 24





EXAMPLE OF SOME TRAFFIC MANAGEMENT CAPABILITIES

FIG. 26